An Interop Publication

CONNEXIONS®

The Interoperability Report

November 1995 Volume 9, No. 11

ConneXions— The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: `connexions@interop.com`

Subscription hotline: 1-800-575-5717
or +1 610-892-1959

ISSN 0894–5926

From the Editor

As mentioned in last month's issue, a textbook on the *Internet Protocol Next Generation* (IPng) has recently been published. The book is a collection of papers by many prominent members of the Internet community. We have obtained permission to publish some selected extracts from this book. Last month we published an article on IPng transition issues. This month we bring you a cable television industry perspective on IPng. The article is by Mario Vecchi of Time Warner Cable and discusses IPng as a potential internetworking technology to support the global requirements of the future integrated broadband networks that the cable industry is designing and deploying.

Asynchronous Transfer Mode (ATM) is perhaps the "hottest" technology of the '90s. In this journal we have previously discussed the standardization work, and have also reported on various ATM testbed efforts. This month we look at some approaches to integrating ATM switching and cell-based routing, and we introduce the *ATM Enterprise Network Switch*. The article is by Tony Rybczynski of Nortel. We have several other ATM-related articles "in the pipeline" for future issues, so stay tuned.

The phenomenal growth of the global Internet is in part the result of careful design and operation of a number of routing protocols. The *Routing Information Protocol* (RIP) has seen widespread use, due in part to the fact that it was bundled with the Berkeley distribution of UNIX (4.x BSD). While the original RIP has been enhanced and respecified by the IETF as RIP Version 2, it still suffers from several limitations making it unsuitable for use in very large internetworks. The *Open Shortest Path First* (OSPF) protocol was developed to address many of the limitations inherent in RIP. In our third article this month, Eddie Rabinovitch of Unisys Corporation discusses migration from RIP to OSPF.

The RISK Internet mailing lists, also known as the USENET newsgroup `comp.risks`, is one of the oldest and most widely read on-line forums. It provides a medium for discussion of issues relating to all aspects of computers and the social and technological problems that they create. The RISKS moderator, Peter Neumann of SRI International, has just published a book entitled *Computer-Related Risks*, and we asked him for a brief essay on computer risks as they relate to communication networks. We are also pleased to announce that Peter will be giving a keynote address at the next NetWorld+Interop in Las Vegas scheduled for April 1996.

A Cable TV Industry View of IPng

by Mario P. Vecchi, Time Warner Cable

Introduction

The IPng requirements and selection criteria from a cable television industry viewpoint are unique. The perspective taken is to position IPng as a potential internetworking technology to support the global requirements of the future integrated broadband networks that the cable industry is designing and deploying. I will include a description of the cable television industry and an outline of the network architectures to support the delivery of entertainment programming and interactive multimedia digital services, as well as telecommunication and data communication services.

Cable networks touch residences, in addition to campuses and business parks. Broadband applications reach the average, computer-shy person. The applications involve a heavy use of video and audio to provide communication, entertainment, and information-access services. The deployment of these capabilities to the homes will represent tens of millions of users. Impact on the network and the IPng requirements that are discussed include issues of scalability, reliability and availability, support for real-time traffic, security and privacy, and operations and network management, among others.

Industry overview

Cable television networks and the Internet are discovering each other. It looks like a great match for a number of reasons, the available bandwidth being the primary driver. Nonetheless, it seems that the impact of the cable television industry in the deployment of broadband networks and services is still not fully appreciated. This article will provide a quick (and simplified) overview of cable television networks, and explain the trends that are driving future network architectures and services.

Cable television networks in the U.S. pass by approximately 90 million homes, and have about 56 million subscribers, of a total of about 94 million homes. There are more than 11,000 head-ends, and the cable TV industry has installed more than 1 million network-miles. Installation of optical fiber proceeds at a brisk pace, the cable fiber plant in the U.S. going from 13,000 miles in 1991 to 23,000 miles in 1992. Construction spending by the cable industry in 1992 was estimated to be about $2.4 billion, of which $1.4 billion was for rebuilds and upgrades. Cable industry revenue from subscriber services in 1992 was estimated to be more than $21 billion, corresponding to an average subscriber rate of about $30 per month (Source: Paul Kagan Associates, Inc.). These figures are based on "conventional" cable television services, and are expected to grow as the cable industry moves into new interactive digital services and telecommunications. [14]

The cable industry's broadband integrated services network architecture is based on a hierarchical deployment of network elements interconnected by broadband fiber optics and coaxial cable links. In a very simplified manner, the following is a view of this architecture. Starting at the home, a coaxial cable tree-and-branch plant provides broadband two-way access to the network. The local access coaxial cable plant is aggregated at a fiber node, which marks the point in the network where fiber optics becomes the broadband transmission medium. Current deployment is for approximately 500 homes passed by the coaxial cable plant for every fiber node, with variations (from as low as 100 to as many as 3000) that depend on the density of homes and the degree of penetration of broadband services.

The multiple links from the fiber nodes reach the head-end, which is where existing cable systems have installed equipment for origination, reception, and distribution of television programming. The head-ends are in buildings that can accommodate weather protection and powering facilities, and hence represent the first natural place in the network where complex switching, routing, and processing equipment can be conveniently located. Traffic from multiple head-ends can be routed over fiber optics to regional hub nodes deeper in the network, where capital-intensive functions can be shared in an efficient way.

Evolution

The cable networks are evolving quite rapidly to become effective two-way digital broadband networks. Cable networks will continue to be asymmetric, and they will continue to deliver analog video. But digital capabilities are being installed very aggressively and a significant upstream bandwidth is rapidly being activated. The deployment of optical fiber deeper into the network is making the shared coaxial plant more effective in carrying broadband traffic in both directions. For instance, with fiber nodes down to where only about 100 to 500 homes are passed by the coaxial drops (down from tens of thousands of homes passed in the past), an upstream bandwidth of several MHz represents a considerable capacity. The recent announcement by Continental Cablevision and PSI of Internet access services is but one example of the many uses of these two-way broadband capabilities.

The cable networks are also rapidly evolving into regional networks. The deployment of fiber optic trunking facilities (many based on SONET) will provide gigabit links that interconnect regional hub nodes in regional networks spanning multiple cable systems. These gigabit networks carry digitized video programming, but will also carry voice (telephone) traffic, and, of course, data traffic. There are instances in various parts of the country where these regional networks have been in successful trials. And given that compressed digital video is the way to deliver future video programs (including interactive video, video on demand, and a whole menu of other applications like computer-supported collaborative work, multiparty remote games, home shopping, customized advertisement, multimedia information services, etc.), one can be guaranteed that gigabit regional networks will be put in place at an accelerated pace.

ATM

The cable networks are evolving to provide broadband networking capabilities in support of a complete suite of communication services. The Orlando network being built by Time Warner is an example of a Full Service Network™ that provides video, audio, and data services to the homes. For the trial, ATM is brought to the homes at DS3 rates, and it is expected to go up to OC-3 rates when switch interfaces are available. This trial in Orlando represents a peek into the way of future cable networks. The Full Service Network uses a "set-top" box in every home to provide the network interface. This set-top box, in addition to some specialized modules for video processing, is a powerful computer in disguise, with a computational power comparable to high-end desktop workstations. The conventional analog cable video channels will be available, but a significant part of the network's RF bandwidth will be devoted to digital services. There are broadband ATM switches in the network (as well as 5E-type switches for telephony), and video servers that include all kinds of movies and information services. An important point to notice is that the architecture of future cable networks maps directly to the way networked computing has developed. General purpose hosts (i.e., the set-top boxes) are interconnected through a broadband network to other hosts and to servers.

continued on next page

Cable TV Industry View of IPng *(continued)*

The deployment of the future broadband information superhighway will require architectures for both the network infrastructure and the service support environment that truly integrate the numerous applications that will be offered to the users. Applications will cover a very wide range of scenarios. Entertainment video delivery will evolve from the current core services of the cable industry to enhanced offerings like interactive video, near video-on-demand and complete video-on-demand functions. Communication services will evolve from the current telephony and low-speed data to include interactive multimedia applications, information access services, distance learning, remote medical diagnostics and evaluations, computer supported collaborative work, multiparty remote games, electronic shopping, etc. In addition to the complexity and diversity of the applications, the future broadband information infrastructure will combine a number of different networks that will have to work in a coherent manner. Not only will the users be connected to different regional networks, but the sources of information—in the many forms that they will take—will also belong to different enterprises and may be located in remote networks. It is important to realize from the start that the two most important attributes of the architecture for the future broadband information superhighway are integration and interoperability. The Internet community has important expertise and technology that could contribute to the definition and development of these future broadband networks.

Engineering considerations

The following comments represent expected requirements of future cable networks, based on the vision of an integrated broadband network that will support a complete suite of interactive video, voice and data services.

Scale

The current common wisdom is that IPng should be able to deal with 10^{12} nodes. Given that there are on the order of 10^8 households in the U.S., we estimate a worldwide number of households of about 100 times as many, giving a total of about 10^{10} global households. This number represents about 1 percent of the 10^{12} nodes, which indicates that there should be enough space left for business, educational, research, government, military, and other nodes connected to the future Internet.

One should be cautious, however, not to underestimate the possibility of multiple addresses that will be used at each node to specify different devices, processes, services, etc. For instance, it is very likely that more than one address will be used at each household for different devices such as the entertainment system (i.e., interactive multimedia "next generation" television(s)), the data system (i.e., the home personal computer(s)), and other new terminal devices that will emerge in the future (such as networked games, PDAs, etc.). Finally, the administration of the address space is important. If there are large blocks of assigned, but unused, addresses, the total number of available addresses will be effectively reduced from the 10^{12} nodes that have been originally considered.

Timescale

The cable industry is already making significant investments in plant upgrades, and the current estimates for the commercial deployment indicate that, by the year 1998, tens of millions of homes will be served by interactive and integrated cable networks and services.

This implies that during 1994, various trials will be conducted and evaluated, and the choices of technologies and products will be well under way by the year 1995. That is to say, critical investment and technological decisions by many of the cable operators, and their partners, will be made by the end of 1995 or 1996.

These time estimates are tentative, of course, and subject to variations depending on economic, technical, and public policy factors. Nonetheless, the definition of the IPng capabilities and the availability of implementations should not be delayed beyond the end of 1995, in order to be ready at the time during which many of the early technological choices for the future deployment of cable networks and services will be made. The full development and deployment of IPng will require, of course, a long period. However, availability of early implementations will allow experimentation in trials to validate IPng choices and to provide early buy-in from the developers of networking products that will support the planned roll out.

The effective support for high-quality video and audio streams is one of the critical capabilities that should be demonstrated by IPng in order to capture the attention of network operators and information providers of interactive broadband services (e.g., cable television industry and partners). The currently accepted view is that IP is a great networking environment for the control side of an interactive broadband system. It is a challenge for IPng to demonstrate that it can be effective in transporting the broadband video and audio data streams, in addition to providing the networking support for the distributed control system.

Transition and deployment

The transition from the current version to IPng has to consider two aspects: support for existing applications and availability of new capabilities. The delivery of digital video and audio programs requires the capability to do broadcasting and selective multicasting efficiently. The interactive applications that the future cable networks will provide will be based on multimedia information streams that will have real-time constraints. That is to say, both the end-to-end delays and the jitter associated with the delivery across the network have to be bounded. In addition, the commercial nature of these large private investments will require enhanced network capabilities for routing choices, resource allocation, quality-of-service controls, security, privacy, etc. Network management will be an increasingly important issue in the future. The extent to which the current IP fails to provide the needed capabilities will provide additional incentive for the transition to occur, since there will be no choice but to use IPng in future applications.

It is very important, however, to maintain backwards compatibility with the current IP. There is the obvious argument that the installed technological base developed around IP cannot be neglected under any reasonable evolution scenario. But in addition, one has to keep in mind that a global Internet will be composed of many interconnected heterogeneous networks, and that not all subnetworks, or user communities, will provide the full suite of interactive multimedia services. Internetworking between IPng and IP will have to continue for a very long time in the future.

Secure operation

The security needed in future networks falls into two general categories: protection of the users and protection of the network resources. The users of the future global Internet will include many communities that will likely expect a higher level of security than is currently available.

continued on next page

Cable TV Industry View of IPng *(continued)*

These users include business, government, research, and military, as well as private subscribers. The protection of the users' privacy is likely to become a hot issue as new commercial services are rolled out. The possibility of illicitly monitoring traffic patterns by looking at the headers in IPng packets, for instance, could be disturbing to most users that subscribe to new information and entertainment services.

The network operators and the information providers will also expect effective protection of their resources. One would expect that most of the security will be dealt with at higher levels than IPng, but some issues might have to be considered in defining IPng as well. One issue relates, again, to the possibility of illicitly monitoring addresses and traffic patterns by looking at the IPng packet headers. Another issue of importance will be the capability of effective network management under the presence of benign or malicious bugs, especially if both source routing and resource reservation functionality is made available.

Configuration, administration, and operation

The operations of these future integrated broadband networks will indeed become more difficult, and not only because the networks themselves will be larger and more complex, but also because of the number and diversity of applications running on or through the networks. It is expected that most of the issues that need to be addressed for effective operations support systems will belong to higher layers than IPng, but some aspects should be considered when defining IPng.

The area where IPng would have most impact would be in the interrelated issues of resource reservation, source routing and quality of service control. There will be a push to maintain high quality of service and low network resource usage simultaneously, especially if the users can specify preferred routes through the network. Useful capabilities at the IPng level would enable the network operator, or the user, to effectively monitor and direct traffic in order to meet quality and cost parameters. Similarly, it will be important to dynamically reconfigure the connectivity among end points or the location of specific processes (e.g., to support mobile computing terminals), and the design of IPng should either support, or at least not get in the way of, this capability. Under normal conditions, one would expect that resources for the new routing will be established before the old route is released in order to minimize service interruption. In cases where reconfiguration is in response to abnormal (i.e., failure) conditions, then one would expect longer interruptions in the service, or even loss of service.

The need to support heterogeneous, multiple administrative domains will also have important implications on the addressing schemes that IPng should support. It will be both a technical and a business issue to have effective means to address nodes, processes, and users, as well as choosing schemes based on fair and open processes for allocation and administration of the address space.

Support for mobility

The proliferation of personal and mobile communication services is a well-established trend by now. Similarly, mobile computing devices are being introduced to the market at an accelerated pace. It would not be wise to disregard the issue of host mobility when evaluating proposals for IPng. Mobility will have impact on network addressing and routing, adaptive resource reservation, security, and privacy, among other issues.

Flows and resource reservation

The largest fraction of the future broadband traffic will be due to real-time voice and video streams. It will be necessary to provide performance bounds for bandwidth, jitter, latency, and loss parameters, as well as synchronization between media streams related by an application in a given session. In addition, there will be alternative network providers that will compete for the users and that will provide connectivity to a given choice of many available service providers. There is no question that IPng, if it aims to be a general protocol useful for interactive multimedia applications, will need to support some form of resource reservation or flows.

Two aspects of flow and resource reservation are worth mentioning. First, the quality of service (QoS) parameters are not known ahead of time, and hence the network will have to include flexible capabilities for defining these parameters. For instance, MPEG-2 packetized video might have to be described differently than G.721 PCM packetized voice, although both data streams represent real-time traffic channels. In some cases, it might be appropriate to provide soft guarantees in the quality parameters, whereas in other cases hard guarantees might be required. The trade-off between cost and quality could be an important capability of future IPng-based networks, but much work needs to be advanced on this.

A second important issue related to resource reservations is the need to deal with broken or lost end-to-end state information. In traditional circuit-switched networks, a considerable effort is expended by the intelligence of the switching system to detect and recover resources that have been lost due to misallocation. Future IPng networks will provide resource reservation capabilities by distributing the state information of a given session into multiple nodes of the network. A significant effort will be needed to find effective methods to maintain consistency and recover from errors in such a distributed environment. For example, keep-alive messages to each node where a queuing policy change has been made to establish the flow could be a strategy to make sure that network resources do not remain stuck in some corrupted session state. One should be careful, however, not to assume that complex distributed algorithms can be made robust by using timeouts. This is a problem that might require innovation beyond the reuse of existing solutions.

It should be noted that some aspects of the requirements for recoverability are less stringent in this networking environment than in traditional distributed data processing systems. In most cases it is not needed (or even desirable) to recover the exact session state after failures, but only to guarantee that the system returns to some safe state. The goal would be to guarantee that no network resource is reserved that has not been correctly assigned to a valid session. The more stringent requirement of returning to old session state is not meaningful since the value of a session disappears, in most cases, as time progresses. One should keep in mind, however, that administrative and management state, such as usage measurement, is subject to the same conventional requirements of recoverability that database systems currently offer.

Policy-based routing

In future broadband networks, there will be multiple network operators and information providers competing for customers and network traffic. An important capability of IPng will be to specify, at the source, the specific network for the traffic to follow. The users will be able to select specific networks that provide performance, feature or cost advantages.

continued on next page

Cable TV Industry View of IPng *(continued)*

From the user's perspective, source routing is a feature that would enable a wider selection of network access options, enhancing their ability to obtain features, performance, or cost advantages. From the network operator and service provider perspective, source routing would enable the offering of targeted bundled services that will cater to specific users and achieve some degree of customer lock-in. The information providers will be able to optimize the placement and distribution of their servers, based on either point-to-point streams or on multicasting to selected subgroups. The ability of IPng to dynamically specify the network routing would be an attractive feature that will facilitate the flexible offering of network services.

Topological flexibility

It is hard to predict what the topology of the future Internet will be. The current model developed in response to a specific set of technological drivers, as well as an open administrative process reflecting the non-commercial nature of the sector. The future Internet will continue to integrate multiple administrative domains that will be deployed by a variety of network operators.

It is likely that there will be more "gateway" nodes (at the head-ends or even at the fiber nodes, for instance) as local and regional broadband networks will provide connectivity for their users to the global Internet.

Applicability

The future broadband networks that will be deployed, by both the cable industry and other companies, will integrate a diversity of applications. The strategies of the cable industry are to reach the homes, as well as schools, business, government, and other campuses. The applications will focus on entertainment, remote education, telecommuting, medical, community services, news delivery, and the whole spectrum of future information networking services. The traffic carried by the broadband networks will be dominated by real-time video and audio streams, even though there will also be an important component of traffic associated with non-time-critical services such as messaging, file transfers, remote computing, etc. The value of IPng will be measured as a general internetworking technology for all these classes of applications. The future market for IPng could be much wider and larger than the current market for IP, provided that the capabilities to support these diverse interactive multimedia applications are available.

It is difficult to predict how pervasive the use of IPng and its related technologies might be in future broadband networks. There will be extensive deployment of distributed computing capabilities, both for the user applications and for the network management and operation support systems that will be required. This is the area where IPng could find a firm stronghold, especially as it can leverage on the extensive IP technology available. The extension of IPng to support video and audio real-time applications, with the required performance, quality, and cost to be competitive, remains a question to be answered.

Datagram service

The "best-effort," hop-by-hop paradigm of the existing IP service will have to be reexamined if IPng is to provide capabilities for resource reservation or flows. The datagram paradigm could still be the basic service provided by IPng for many applications, but careful thought should be given to the need to support real-time traffic with (soft and/or hard) quality of service requirements.

Accounting

The ability to do accounting should be an important consideration in the selection of IPng. The future broadband networks will be commercially motivated, and measurement of resource usage by the various users will be required. The actual billing may or may not be based on session-by-session usage, and accounting will have many other useful purposes besides billing. The efficient operation of networks depends on maintaining availability and performance goals, including both on-line actions and long-term planning and design. Accounting information will be important on both scores. On the other hand, the choice of providing accounting capabilities at the IPng level should be examined with a general criterion to introduce as little overhead as possible. Since fields for "to," "from," and time stamp will be available for any IPng choice, what other parameters in IPng could be useful to both accounting and other network functions should be examined, so as to keep IPng as lean as possible.

Media independence

The generality of IP should be carried over to IPng. It would not be an advantage to design a general internetworking technology that cannot be supported over as wide a class of communications media as possible. It is reasonable to expect that IPng will start with support over a few select transport technologies, and rely on the backwards compatibility with IP to work through a transition period. Ultimately, however, one would expect IPng to be carried over any available communications medium.

Robust service

Service availability, end-to-end, and at expected performance levels, is the true measure of robustness and fault-tolerance. In this sense, IPng is but one piece of a complex puzzle. There are, however, some vulnerability aspects of IPng that could decrease robustness. One general class of bugs will be associated with the change itself, regardless of any possible enhancement in capabilities. The design, implementation, and testing process will have to be managed very carefully. Networks and distributed systems are tricky. There are plenty of horror stories from the Internet community itself to make us cautious, not to mention the brief but dramatic outages over the last couple of years associated with relatively small software bugs in the control networks (i.e., CCS/SS7 signaling) of the telephone industry, both local and long distance.

A second general class of bugs will be associated with the implementation of new capabilities. IPng will likely support a whole set of new functions, such as larger (multiple?) address space(s), source routing, and flows, just to mention a few. Providing these new capabilities will require in most cases designing new distributed algorithms and testing implementation parameters very carefully. In addition, the future Internet will be even larger, have more diverse applications and have higher bandwidth. These are all factors that could have a multiplying effect on bugs that in the current network might be easily contained. The designers and implementors of IPng should be careful. It will be very important to provide the best possible transition process from IP to IPng. The need to maintain robustness and fault-tolerance is paramount.

Technology pull

The strongest "technology pull" factors that will influence the Internet are the same that are dictating the accelerated pace of the cable, telephone, and computer networking world.

continued on next page

Cable TV Industry View of IPng *(continued)*

The following is a partial list: higher network bandwidth, more powerful CPUs, larger and faster (static and dynamic) memory, improved signal processing and compression methods, advanced distributed computing technologies, open and extensible network operating systems, large distributed database management and directory systems, high-performance and high-capacity real-time servers, friendly graphical user interfaces, and efficient application development environments. These technology developments, coupled with the current aggressive business strategies in our industry and favorable public policies, are powerful forces that will clearly have an impact on the evolution and acceptance of IPng. The current deployment strategies of the cable industry and their partners do not rely on the existence of commercial IPng capabilities, but the availability of new effective networking technology could become a unifying force to facilitate the interworking of networks and services.

Summary

The potential for IPng to provide a universal internetworking solution is a very attractive possibility, but there are many hurdles to be overcome. The general acceptance of IPng for supporting future broadband services will depend on more than the IPng itself. There is need for IPng to be backed by the whole suite of Internet technology that will support the future networks and applications. These technologies must include the adequate support for commercial operation of a global Internet that will be built, financed, and administered by many different private and public organizations.

The Internet community has taken pride in following a nimble and efficient path in the development and deployment of network technology. And the Internet has been very successful up to now. The challenge is to show that the Internet model can be a preferred technical solution for the future. Broadband networks and services will become widely available in a relatively short future, and this puts the Internet community in a fast track race. The current process to define IPng can be seen as a test of the ability of the Internet to evolve from its initial development—very successful but also protected and limited in scope—to a general technology for the support of a commercially viable broadband marketplace. If the IPng model is to become the preferred general solution for broadband networking, the current IPng process seems to be a critical starting point.

References

[1] Bradner, S., and Mankin, A., "IP: Next Generation (IPng) White Paper Solicitation," RFC 1550, December 1993.

[2] Braden, R. Clark, D., Shenker, S., "Integrated Services in the Internet Architecture: An Overview," RFC 1633, June 1994.

[3] Clark, D., "The Design Philosophy of the DARPA Internet Protocols," Proceedings of SIGCOMM '88, *Computer Communications Review,* Volume 18, No. 4, August 1988.

[4] The ATM Forum, "ATM User–Network Interface, Version 3," September 1993.

[5] "Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements," Bellcore Technical Advisory TA-NWT-001113, Issue 1, June 1993.

[6] Droms, R., "Dynamic Host Configuration Protocol," RFC 1541, October 1993.

[7] Rekhter, Y., and Li, T., "An Architecture for IP Address Allocation with CIDR," RFC 1518, September 1993.

[8] HPN Planning Group, "Concepts and Guidance for High Peformance Network (HPN)," Work in Progress, May 1993.

[9] Rivkin, S., "Positioning the Electric Utility to Build Information Infrastructure," U.S. Department of Energy.

[10] Cass, D. and Rose, M. T., "ISO Transport Services on Top of the TCP: Version: 3," RFC 1006, May 1987.

[11] Institute for Simulation and Training, "Communication Architecture for Distributed Interactive Simulation (CADIS)," June 1993.

[12] Miller, D., "Distributed Interactive Networking Issues," Briefing presented to the ST/IP Peer Review Panel, Massachusetts Institute of Technology, Lincoln Laboratory, December 1993.

[13] Braden, R., Zhang, L., Estrin, D., Herzog, S., and Jamin, S., "Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification," Work in Progress, January 1995.

[14] United States Television Census, September 1993.

[15] Hinden, R. "IP Next Generation Overview," *ConneXions,* Volume 9, No. 3, March 1995.

[16] *ConneXions,* Volume 8, No. 5, May 1994, Special Issue on IP The Next Generation.

[17] Gilligan, R., and Callon, R., "IPv6 Transition Mechanisms Overview," *ConneXions,* Volume 9, No. 10, October 1995.

[18] Bradner, S., & Mankin, A., "The Recommendation for the IP Next Generation Protocol," RFC 1752, January 1995.

[19] Hinden, R., Editor, "IP Version 6 Addressing Architecture," Work in Progress, April 1995.

[20] Rekhter, Y., & Li, T., "An architecture for IPv6 Unicast Address Allocation," Work in Progress, March 1995.

[21] Rekhter, Y., & Löthberg, P., "An IPv6 Global Unicast Address Format," Work in Progress, March 1995.

[22] Fleischman, E., "A User's View of the Next Generation of IP (IPng)," *ConneXions,* Volume 8, No. 5, May 1994.

MARIO P. VECCHI is currently Vice President, Network Engineering at Time Warner Cable, Engineering and Technology, Englewood, CO, where he is working to develop the technology and products for delivery of high-speed data services to home personal computers over the broadband hybrid fiber/coax network using multi-megabit/s cable modems. He is interested not only in the deployment of high-speed data applications to residential PCs, but also in the integration with telecommunications and interactive video entertainment, based on open, scalable, and extensible distributed systems architectures. E-mail: `Mario.Vecchi@twcable.com`

[Ed.: This article is adapted from the book *IPng: Internet Protocol Next Generation,* Edited by Scott O. Bradner and Allison Mankin, published by Addison-Wesley, ISBN 0-201-63395-7, 1995. Used with permission].

The ATM Enterprise Network Switch:
Enabler of Network Transformation

by Tony Rybczynski, Nortel

Introduction

A revolution is in progress that is dramatically changing the way information networks are designed. The introduction of a new form of switching, *Asynchronous Transfer Mode* (ATM), in both in-building and wide area environments will revolutionize the way networks of the future operate. ATM switching will extend LAN application performance across the wide area, making network capacity more scalable, simplifying network management and improving bandwidth price/performance (through the consolidation of other traffic including voice and video). Network transformation is the process of evolving corporate networks towards multimedia, dynamic bandwidth, ATM-based infrastructures. The benefits being sought by users from this network transformation are:

- Improved performance for a broad range of applications
- Greater scalability from multi-Mbit/s to multi-Gbit/s from work group to wide area
- Simplification via multimedia networking consolidation and via simplified management inherent in connection-oriented environments.

While they are increasingly establishing ATM as their long term direction, users are being driven now to simplify their operations environments, manage their capital investments, and minimize their recurring costs, while meeting the needs of a diverse set of demanding end users.

In the local environment, this is primarily being accomplished by the movement beyond shared-media LANs to switched Ethernet/Token Ring solutions. Ethernet and Token Ring switches based on ATM (and therefore including ATM LAN emulation) provide an evolution path to ATM. The introduction of LAN switching enhances application performance and traffic scalability, while decreasing the administrative complexity of shared-media LANs and router configurations. In cases where 10/16 Mbit/s is insufficient, there are numerous desktop technologies such as 100 Mbit/s Fast Ethernet, 100 Mbit/s VG AnyLAN, FDDI and ATM. The winning technology will be the one that delivers the best price/performance, while minimizing operational impacts.

In the wide area environment, many users have deployed TDM multiplexor backbones. TDM multiplexors support channelized operation and are very ineffective for bursty inter-router traffic, are typically limited to T1 speeds and do not support ATM interfaces and switching. In more recent years, there has been a trend to application specific networks with router backbones (increasingly interconnected via Frame Relay) for inter-LAN connectivity, virtual private voice networks, and video networks, and now emerging ATM networks. ATM offers the opportunity to consolidate these into a single network.

An important objective is to build networks that maximize application performance and enhance service predictability through the support of various classes of service. With ATM switching exhibiting latencies of at least a hundred times better than routers, it seems obvious that latency can be minimized by minimizing routing in favor of switching.

This article will discuss numerous approaches to integrating ATM switching and cell-based routing, and introduce a new wide area product class which has evolved in the marketplace: the *ATM Enterprise Network Switch* (ATM ENS), which consolidates cell-based routing, ATM switching and multimedia consolidation on a single high performance platform.

We will then discuss various approaches to consolidating voice and video traffic over an ATM ENS, with a particular focus on traffic management implications. This is critical to allow users to achieve the benefits of network consolidation of multimedia traffic. Comprehensive traffic management is a key requirement for ATM-based networks. While *Unspecified Bit Rate* (UBR) virtual circuits were initially used for inter-LAN connectivity, the need to define mechanisms to adjust the rate of traffic to network congestion conditions was identified. This resulted in some ATM NIC card vendors defining proprietary end-to-end closed loop congestion control mechanisms. It also resulted in the ATM Forum defining rate-based flow control mechanisms in the *Available Bit Rate* (ABR) class of service standard. An initial common approach to consolidating voice and video over ATM networks was to use constant bit rate *Permanent Virtual Circuits* (PVCs) through what is known as *AAL1 Adaptation*. For the same reason that drove the definition of ABR including dynamic rate adaptation at the traffic source, real-time variable bit rate approaches provide application specific ways of dealing with network congestion, this being done by allowing the coding rate for voice and video to vary in response to network congestion conditions. Voice is discussed in some detail in this article because of its prevalence, while video is discussed at a relatively higher level.

The article will then discuss two key transition issues in transforming corporate networks through the deployment of ATM ENS based networks:

- The need to support migration from virtual private voice networks used by many corporations to ATM ENS-based networks through *Switched Virtual Circuit* (SVC) support. SVCs can also enhance private network operation.
- The need to allow connectivity to remote sites using more bandwidth effective alternatives to ATM transmission, when the use of ATM would not be economically viable (e.g., data only or multimedia connectivity over expensive long haul, typically T1 and fractional T1, facilities).

Cell-based routing and ATM switching

The key attributes of ATM for Inter-LAN connectivity are as follows:

- High performance and low latency, low overhead switched virtual circuit connectivity with the ability to burst at line rate
- ATM classes of service on a per connection basis to enable simultaneous support for different applications
- Scalable capacity with ATM *User Network Interfaces* (UNIs) at speeds up to 600 Mbit/s (OC-12c)
- Low loss through end-to-end signaling with evolution to rate-based flow control based on emerging ABR standards
- Management simplification through connection-oriented operation

continued on next page

The ATM Enterprise Network Switch *(continued)*

A number of approaches can be used to combine multiprotocol routing and ATM switching, to provide inter-LAN connectivity in the wide area, while leveraging the above attributes of ATM. Routers can be interconnected over an ATM network, but this does not eliminate the latency introduced by routers. This approach provides loose coupling between routing and ATM connection and traffic management.

A second approach has been proposed by a number of vendors and relies on routing functionality being distributed between centralized route servers and work group level packet forwarders. This approach establishes a single point of failure and a potential bottleneck in the route server and could have a negative impact on network latency (e.g., by introducing extra hops).

A third approach calls for a single device, such as a cell-based router or an ATM Enterprise Network Switch, to integrate conventional routing and ATM switching. ATM switching provides high-speed connectivity for IP encapsulated on ATM (i.e., classic IP), and for ATM-connected workstations and switched Ethernet/Token Ring hubs using ATM LAN Emulation. ATM switching also provides high speed connectivity for emerging native ATM applications. Integrated multiprotocol routing supports communication with the installed base of shared-media LAN bridge/router networks, between emulated ATM LANs and manages broadcasts and unknown address flooding over the wide area. This approach offers the lowest application delay by choosing optimal switched paths and by eliminating multiple networking devices between end systems. It integrates routing, address resolution, packet forwarding and ATM connection, traffic, and quality of service management. Furthermore, it can provide "circuits" for the network consolidation of voice, data and video.

An ATM ENS provides a range of in-building and wide-area ATM UNIs and native ATM switching. *Constant Bit Rate* (CBR), *Variable Bit Rate* (both real-time and non-real-time VBR), *Unspecified Bit Rate* (UBR) and *Available Bit Rate* (ABR) classes of service are supported to provide more predictable application performance and to optimize the use of network resources. As public network ATM services become available with competitive price/performance, public network ATM interfaces can be added as required. These will support ATM ENS-to-ATM ENS communications, as well as connectivity to routers and ATM multiplexors that support standard ATM Adaptations (i.e., AAL1 for circuit emulation, AAL5 for data).

LANE

ATM-based switching hubs and ATM workgroup switches are being deployed by end users at various corporate sites and will require wide area connectivity. The most widely used set of standards in local ATM environments is *ATM LAN Emulation* (LANE). ATM LAN Emulation is used to make the ATM network appear to be a collection of virtual Ethernet/IEEE 802.3 and Token Ring/IEEE 802.5 LANs. The replication of most of the characteristics of existing LANs means that LAN Emulation enables existing LAN applications to run over ATM transparently, this latter characteristic leading to its wide deployment. In LANE, most unicast LAN traffic moves directly between clients over direct ATM VCs, while multicast traffic is handled via a server functionality. Bridging is used to interconnect real LANs and emulated LANs running on ATM, while routing is used to interconnect ATM emulated LANs and other wide area or LAN media for purposes of routing scalability, protocol spoofing, or security firewalls.

The ATM Forum LANE implementation agreement specifies two types of LANE network components connected to an ATM network:

- LANE clients are end systems such as:
  - Computers with ATM interfaces that operate as file servers
  - End-user workstations or personal computers
  - Ethernet or Token Ring switches that support ATM networking.
  - Routers, bridges and ATM ENSs with membership in an emulated ATM LAN
- LANE servers supporting ATM LANE Service for configuration management, multicast support and address resolution.

An ATM ENS interfaces to ATM workgroup switches and ATM-based LAN switches using in-building ATM UNIs (e.g., based on multimode fiber), appears as a LANE client on an emulated ATM LAN, and provides inter-emulated LAN connectivity via routing for the most common protocols. Emulated LANs can be extended across the wide area using ATM virtual circuits to establish a distributed virtual LANs. An ATM ENS also supports native ATM switching to support new multimedia applications that fully exploit the characteristics of ATM.

Over time, an ATM ENS will have to participate in the *Multiple Protocol over ATM* (MPOA) environment, that is being defined by the ATM Forum. This work is addressing encapsulation of multiple protocols over ATM, automatic address resolution, and the routing issues associated with minimizing multiple router hops in ATM networks. An ATM ENS could provide distributed networking server functionality for the optimum combination of routing and ATM switching over a wide area network.

Voice and video consolidation

Numerous voice over ATM applications exist. In the public network, transporting voice calls originated from analog or digital telephones—including mobile cellular users over an ATM backbone network is an opportunity being assessed by carrier network planners. In enterprise networks, the requirement is to support inter-PBX trunks over ATM, with ATM adaptation done by the ATM ENS. In the longer term, ATM ENS networks will also have to support voice communications which is adapted to ATM right at the desktop. Two cases can be envisaged:

- Voice supported at the desktop as part of an ATM-specific collaborative application, and
- As a replacement for the conventional telephone for business communications.

In the latter case, feature transparency must be provided across the ATM and conventional PBX environments (e.g., call forwarding, ring again, three way calling). Here, we focus on inter-PBX transport across an ATM ENS network. While voice traffic, may not be growing significantly in corporate networks, voice represents a significant percentage of the communications bill each month. The saving of potentially many cents per minute of voice connect time is the major driver behind consolidating PBX voice trunking over ATM networks.

continued on next page

The ATM Enterprise Network Switch *(continued)*

Two distinct networking approaches for PBX networking have been widely deployed, specifically (i) private voice networks; and (ii) virtual private voice networks. In private voice networks, the user establishes his own networking topology using private lines or TDM multiplexor-derived channels. Direct routes are established across major cross-sections while tandem PBXs are defined to provide connectivity among all sites. Uniform dialing and feature transparency are supported by inter-PBX signaling systems and by administering network routing tables in each PBX. In virtual private voice networks, each PBX is connected to the public network, which interprets the signaling received to route the call to the destination PBX. In this case the routing tables are managed on a centralized basis as part of the carrier service. Even when virtual private voice networks are implemented, direct traffic between major corporate sites can be handled over private networks without complicating routing table administration. Therefore, this section will address the adaptation of PBX voice trunks over an ATM ENS network using ATM PVCs.

An initial common approach to consolidating voice trunks over ATM networks was to use constant bit rate PVCs and AAL1 adaptation, this being done at the T1 level. In this case, 1.544 Mbit/s (24 DS0 channels in a T1) is conveyed across the network as a 1.74 Mbit/s ATM CBR PVC. This approach is very ineffective as only some of the 24 channels are defined as PBX trunks, and only a small percentage of these are actually in use at any given time. A more effective approach is to operate only on individual channels and map each predefined DS0 into its own CBR PVC. This also has the advantage that each DS0 can be routed independently. With DS0, and in some cases T1, circuit emulation, echo cancellation will be required to meet voice performance requirements.

Benefits

There are a number of advantages to using VBR-real-time VCs. VBR-real time VCs would recognize the inherently bursty nature of voice communication, as there are silence periods which can result in increased efficiency. These silence periods (in decreasing levels of importance) arise:

- When no call is up on a particular trunk; that is, the trunk is idle during off-peak hours (trunks are typically engineered for a certain call blocking probability: at night, all the trunks could be idle)
- When the call is up, but only one person is talking at a given time
- When the call is up, and no one is talking

The addition of more bandwidth effective voice coding (e.g., standard voice is coded using 64kbit/s *Analog to Digital Pulse Code Modulation* [ADPCM]) is economically attractive particularly over long-haul circuits and T1 ATM interfaces. Making these coding schemes dynamic provides the network operator the opportunity to free up bandwidth under network congestion conditions. For example, with the onset of congestion, increased levels of voice compression could be dynamically invoked, thus freeing up bandwidth and potentially alleviating the congestion condition, while diminishing the quality of the voice during these periods. A further enhancement to the support of voice over an ATM ENS network is to support voice switching over SVCs. This entails interpreting PBX signaling and to route voice calls to the appropriate destination PBX.

This will be discussed in the next section. The advantage from a traffic management perspective is that connection admission controls can be applied to new voice calls; under network congestion conditions, these calls could be rerouted over the public network, and therefore not cause additional levels of congestion.

Turning briefly to video, there are numerous application areas including room-to-room video conferencing (e.g., at $n \times$ 64kbit/s or multi-Mbit/s), broadcast quality TV (e.g., for training), desktop video over Ethernet, and residential video-on-demand. An important area being addressed is the use of ATM to support higher quality video in the $n \times$ T1 range (using MPEG2 coding). ATM is well suited to video because it can provide the bandwidth required for a given video quality, recognizing the rapid evolution of video codec technology. For example, video codecs typically allow user control over the bit rate for the desired video quality. The initial approach is to consolidate video circuits over ATM networks using constant bit rate PVCs and AAL1 or AAL5 (for MPEG2) adaptation, this being done at the $n \times$ T1 rates.

However, once again, there are a number of advantages to using VBR-real-time VCs. VBR-real time VCs could exploit the inherently bursty nature of video communication, as there are periods of relative inaction. The MPEG2 video coding standard, which is being targeted for ATM, in fact leverages these characteristics of video. Making these coding schemes dynamic provides the network operator the opportunity to free up bandwidth under network congestion conditions. For example, with the onset of congestion, increased levels of video compression could be dynamically invoked, thus freeing up bandwidth and potentially alleviating the congestion condition, while diminishing the quality of the video during these periods.

Migration issues

In a typical PBX implementation, call routes are preconfigured and static in each switch. Primary and alternate routes are specified per destination taking into account peak traffic bandwidth required on each route for a particular call blocking probability. Pre-configured hop-by-hop routes may result in blocked calls in complex network topologies. In such cases, calls may be blocked at that point or overflowed to the public network. Addition of a new PBX alters the routing table in each PBX, while bandwidth in the new network needs be recomputed. Addition of a new prefix alters the routing table in each PBX. Operations costs are typically high due to network design/redesign as traffic patterns change, and increase exponentially with the size of the network.

While some customers have achieved significant economic advantage by carrying voice over TDM multiplexor networks and are continuously looking for ways to simplify the administration of these networks, others have gone with virtual private voice networks. Virtual private voice network services have been introduced by service providers over the years providing economic advantages and simplified administration, by eliminating the need for complex routing table administration at each PBX. Virtual private voice network customers will only consider moving their voice networking needs onto an ATM enterprise network if there are demonstrated economic incentives (there generally are) and if the simplicity of administering PBX routing is maintained.

Support for voice switching over ATM supports the needs of both types of users. Each PBX is connected to an ATM ENS, which interprets the signaling received to route the call directly to the destination PBX.

continued on next page

The ATM Enterprise Network Switch *(continued)*

The number of PBX interfaces is determined by the total peak load, not by the number of trunk groups to remote PBXs. Voice routing tables are managed on a centralized basis as is the case with virtual private voice networks. Routes are calculated dynamically on a per call basis, based on the current network conditions thus simplifying network design. A single update for the whole network (not node by node) is sufficient to add or remove a PBX in the network.

The number of telecommuters and mobile users is growing rapidly. Application of ATM ENS voice networking could also be used to allow telecommuters and mobile users to access the corporate voice network, by making the public voice network appear logical as a PBX interfacing the ATM ENS. This would require the addition of calling line identification for user authentication. ATM ENS voice networking based on VBR-real-time SVCs provides network economics and low overhead voice network routing table administration.

Non-ATM operation for effective bandwidth use

The cost of wide area bandwidth is the dominant cost in enterprise networks, in many cases accounting for over 80% of lifecycle network costs. In addition, enterprise networks typically consist of major and branch domestic sites, and international locations. The challenges faced by network planners within tight budgetary constraints are:

- To meet application needs
- To extend the network coverage
- To maximize performance
- To manage the network bandwidth under congestion and failure conditions.

The use of ATM as the only method of interconnecting sites results in either bandwidth inefficiencies or the continued use of non-ATM vehicles at more remote, international or smaller sites.

The process of adapting data frames to ATM cells results in overheads of as high as 40% (e.g., typical TCP/IP packets of 64 bytes). On the other hand, the same packet sent in its native mode would incur an overhead of only a few percent. Sending long frames can bring the overheads below a few percentage points. However, sending long frames without segmentation has the impact that delay variations become excessive for real-time traffic. ATM ENS architectures, that incorporate hardware support for both standard ATM-based networking and frame/cell trunking, provide network operators with the option of increased bandwidth efficiency on certain cross-sections.

Two approaches have been developed to deal with the delay variation issue associated with supporting real-time traffic and frame data on a single interface. One approach uses variable length cells (up to some maximum to meet delay variation objectives), allowing the user to select maximum cell sizes to meet performance needs. In another approach data frames are interrupted to support a short burst (or cell) of isochronous (i.e., voice or video) information; the remainder of the frame is sent once the cell has been transmitted. This latter approach is more efficient than the former because delay variation requirements can be met while transmitting longer data frames.

In addition, in both ATM and frame cell mode, further bandwidth efficiencies are achieved by using variable bit rate voice and video ATM adaptation (as discussed in the previous section). Additionally, variable bit rate techniques can also be used for data circuits by suppressing idle frames (e.g., flags for HDLC frame-based data).

Support for non-ATM operation allows users to start the transformation of their wide area network towards multimedia, dynamic bandwidth, ATM-based infrastructures, by deploying an ATM-based switch architecture and adding ATM interfaces (to support ATM workgroup switches or to use carrier ATM services) on a plug and play basis.

Conclusions

Network transformation is the process of evolving corporate networks towards multimedia, dynamic bandwidth, ATM-based infrastructures. An ATM Enterprise Network Switch is a new class of product that has evolved as the optimal convergence of routing and switching (supporting cell-based routing) and as a multimedia network consolidation vehicle. While pure AAL1 circuit emulation has its role, support of voice and video over VBR-real-time VCs allows the network to adapt the adaptation rate to network congestion conditions, while achieving improved network economics. Support for voice over ATM SVCs can ease the operation of voice networks, and ease the transition from virtual private voice networks while further improving network economics. In order to allow evolution to ATM public services on a plug and play basis, and to start the transformation of the corporate network to ATM-based multimedia dynamic bandwidth networking, some ATM ENS architectures also support integrated frame/cell trunking options.

The major challenges associated with LAN networking are performance, scalability, management and cost. In the local environment, users are rapidly moving to switched LAN architectures and ATM. In response to user requirements for LAN interconnectivity across the campus or over an international wide area network, various vendors have developed ATM Enterprise Network Switches, supporting cell-based routing and sophisticated ATM adaptation capabilities for voice and video traffic. These ATM ENSs deliver the benefits of switching architectures across the wide area, making network capacity more scalable, simplifying the network management and though traffic consolidation including voice and video, and improving bandwidth price performance. These are effective vehicles to support business case driven enterprise network transformation to ATM-based networking.

References

[1] Laubach, Mark, "ATM for your internet—But When?" *ConneXions,* Volume 7, No. 9, September, 1993.

[2] Laubach, Mark, "Classical IP and ARP over ATM," RFC 1577, January, 1994.

[3] Atkinson, Ran, "Towards Real ATM Interoperability," *ConneXions,* Volume 7, No. 8, August 1993.

[4] Laubach, M., "IP Over ATM and the Construction of High-Speed Subnet Backbones," *ConneXions,* Volume 8, No. 7, July, 1994.

[5] Clapp, G. & Zeug, M., "Components of OSI: Asynchronous Transfer Mode," *ConneXions,* Volume 6, No. 4, April 1992.

TONY RYBCZYNSKI is Consulting Director, Strategic Marketing, in Nortel's Multimedia Networks Group, located in Ottawa. He has worked in the industry for 23 years. After graduating from McGill and University of Alberta, he joined the Computer Communications Group (a business unit of Telecom Canada—now Stentor—operating within Bell Canada) and became one of the "Fathers of X.25" and was a founding member of Datapac. In 1982, he joined Bell Northern Research and was responsible for strategic product planning. While at BNR, he lead the development of a Frame Relay testbed. Since 1986, he has been with NORTEL Magellan Networks. He has been instrumental in the definition of the Magellan product line, and has recently spent considerable time working with end users on in-building and wide area applications of ATM. He has delivered over 50 papers at various conferences worldwide and co-authored a reference book on computer communications protocols. E-mail: `tony.rybczynski@nt.com`

Migration to OSPF is not a Luxury

by Eddie Rabinovitch, Unisys Corporation

Introduction

Any IP address consists of an network part and local part. The length of the network part depends on the Internet network Class assigned by The IANA registry: e.g., IBM is a class A network and its Internet part is bitwise: "`0000 1001`" (`x'09'`). To allow subnetworking in this network, a technique based on subnet masking is used. This is an indication, known to all routers in this network, that tells which part of the local part of the IP address is to be considered as a subnetwork portion and which part is the host portion in that subnetwork. With subnet masking the semantics of the IP address takes the form:

Network part			Host part	
9. Octet 1	67. Octet 2	38. Octet 3	76 Octet 4	IP Address
0000 1001	0100 0011	0010 0110	0100 1100	
255.	255.	255.	192	Subnet Mask
1111 1111	1111 1111	1111 1111	1100 0000	

This corresponds to host "12" in network `9.67.38.64`

The IP routing algorithm is straightforward: if the IP part and subnetwork part match a directly connected network, then send the IP datagram to the destination IP address on that subnet (resolving the complete IP address into a physical address (e.g., MAC address using ARP). Otherwise, take the destination IP address and look up the routing table to find the entry that forms the longest match with the destination IP address. If matching entry found, route to the indicated IP address (next hop). If no match found: discard the packet (and send an ICMP "Destination Unreachable" message to the sender).

TCP/IP products can use static routes definitions or dynamic routing protocol: e.g., RIP (*Routing Information Protocol*), OSPF (*Open Shortest Path First*), Border Gateway Protocols, and other dynamic routing schemes, to exchange routing information to resolve routing errors. This article compares the two most important routing mechanisms for interior (intra-domain) routing: RIP and OSPF.

RIP

RIP, an Interior Gateway Protocol (defined by RFC 1058 [6]), is the most widely accepted routing protocol. It is also known by the name of a program that implements it in UNIX named *RouteD* that was originally designed at the University of California at Berkeley to provide consistent routing and reachability information among machines on a local area network. RIP's popularity is not necessarily based on its technical merits, but it probably resulted because UC Berkeley distributed *RouteD* along with their popular 4.x BSD UNIX systems. Thus, many Internet sites adopted and installed *RouteD* and started using RIP without even considering its technical merits and limitations. Once installed and running, it became the basis for local routing.

The underlying RIP is straightforward: it arranges to have routers broadcast their entire current routing database periodically: typically every 30 seconds. This message lists each destination along with a distance to that destination measured in number of router hops. RIP is also known as a *Distance-Vector* routing algorithm, which means there is a distance, a cost, and a vector for each destination (the vector just shows the name of the neighboring router, not the entire path).

RIP problems

As a distance-vector based algorithm, RIP works fine for small and stable high-speed network. Instead of passing along the status of the links to the networks, the router tells its neighbors about the entire world, where the best link is precomputed and not necessarily the fastest one. And in times of network instability, because every router is broadcasting his entire routing table, it takes a while for those states to converge into a common view of the network topology. Specifically, RIP does not address well three major types of problems:

- A RIP based network has no simple provisions for protection from routing loops within a network, therefore, to a large extent implementations must trust all network participants to prevent such loops during operation.
- RIP uses a hop count of 15 do denote infinity, which makes it unsuitable for large networks.
- RIP creates so called slow convergency or *Count to Infinity* problems in which inconsistencies arise, because routing update message propagate slowly across the network. Particularly in large networks (or networks with slow links) some routers may still advertise a route that has vanished. (That, by the way, was one of the reasons 15 was chosen as the value of infinity to limit the slow convergency effect).

Slow convergency can be addressed by a technique called *Hold Down,* which forces the router to ignore information about a network for a fixed period of time (typically 60 seconds). The idea is to wait long enough that all machines receive the bad news about a vanished link and do not mistakenly accept outdated information. The disadvantage of this technique is that incorrect routes and routing loops will be preserved for the duration of the Hold Down, even when a valid alternate path is available.

Another, more popular approach to address the slow convergency problem is a technique known as *Split Horizon* update. This technique is widely used by router vendors. With Split Horizon, a router records the interface over which it received a particular route and does not propagate its higher cost route back over the same interface. However, Split Horizon does not resolve the slow convergency problem for all topologies. It also introduces a problem for a Frame Relay network that would not permit full-meshed connectivity for partially meshed networks.

RIP allows the use of subnet masks. The subnet masks are static information defined in the routers. RIP has no provision for exchanging subnet mask information between routers. The routers in your network should use the same subnet masks as RIP to make IP datagrams travel from source to destination within your network. This is another technical limitation of RIP. For example, with Frame Relay, the entire network can be assigned one subnet, but as was shown above, because of the Split Horizon technique, RIP would not permit full-meshed connectivity in partially meshed Frame Relay networks.

continued on next page

Migration to OSPF is not a Luxury *(continued)*

If the subnet mask is not constant in a RIP based network, it becomes impossible for the routers to distinguish between the network part and the host part, since RIP cannot dynamically update/change the mask. Hence throughout the subnet one mask should be used. Therefore, if a variable subnet mask is needed, static routes (or more advanced routing protocols—e.g., RIP-2 or OSPF) have to be used.

RIP-2

In 1993 IETF defined RIP-2 (RFC 1388 now RFC 1723 [4]) to address some of the problems inherent in the original RIP-(1). In particular, RIP-2 supports the following new features: *Routing Domains, External Route Tags,* subnet masks, *Next Hop Addresses,* and authentication. Routing domains allow multiple RIP "clouds" to exist over the same physical network. The *External Route Tag* field was defined in RIP-2 to propagate information acquired from an external gateway, for example, it can be used to propagate an EGP *Autonomous System* (AS) number. One of the major problems of the original RIP was its inability to handle variable subnet masks. Inclusion of subnet masks was the most important goal for opening the RIP protocol for improvement. Subnet masks are necessary for implementation of "classless" addresses as defined in CIDR (*Classless Inter-Domain Routing* as specified in RFC 1519), which allows more efficient use of the existing 32-bit address space.

An additional improvement with RIP-2 is the new support for Next Hop Addresses, that allows for optimization of routes in an environment which uses multiple routing protocols. Addition of an authentication mechanism, similar to one included in OSPF makes RIP-2 networks more secure. RIP-2 packets may be multicast instead of being broadcast, which reduces the load on hosts, that do not support routing protocols. It also allows RIP-2 routers to share information, which is not understood (and could confuse) RIP-1 routers. Since there are still significantly more RIP based networks implementations than the newer link-state based IGPs, the new features defined in RIP-2 increased its usefulness and made it a viable option to be considered in certain situations. With the advent of OSPF and IS–IS, many believe that RIP is obsolete. While it is true that the newer IGP routing protocols are far superior to RIP, RIP does have some advantages, primarily, in a small networks. For example, RIP has very little bandwidth overhead and configuration and management time. RIP is also very easy to implement, especially in relation to the newer IGPs, like OSPF or IS–IS. Additionally, there are many, many more RIP implementations in the field than OSPF and IS–IS combined, and it is likely to remain that way for some years yet.

OSPF

OSPF (RFC 1247 now RFC 1583) is a *Link-State Algorithm.* A distance-vector router advertises a set of destinations reachable through the router, while a link-state router advertises network topology expressed as a set of links. Or simply speaking, a distance-vector router "tells all neighbors about the world," while a link-state router "tells the world about the neighbors." OSPF specifies a class of messages called *Link-State Advertisements* (LSAs) that allow routers to tell each other about the LAN/WAN links to which they are connected. When a change is made to the network, LSAs flow between routers. [5]

OSPF routers receive link-state updates and store them in a topology database in memory. The typical OSPF database contains a representation of every link and router in the enterprise network. When routers receive internetwork traffic that needs to be forwarded towards a destination end-node, they use their topology database to calculate a table of the best routes through an internet.

OSPF addresses all problems that exist with RIP and is therefore better suited for modern large and dynamic networks. For example, in contrast to RIP sending the entire routing table from router to router every 30 seconds. OSPF sends its link state information every 30 *minutes*. OSPF can get away with this, because OSPF routers also send each other small update messages (typically less than 75 bytes) whenever they detect a change in the network (e.g., a failure or a new link). When routers exchange updates that reflect changes in the network, they "converge" on a new representation of the topology quickly and accurately.

Variable subnet addressing is one of the notable advantages of the OSPF protocol versus the original RIP. It allows one to maximize the number of viable IP addresses within an organization; in TCP/IP networks, the 32 bits of an IP address are divided up into parts that identify an end-node and its attached network; variable subnet addressing allows the 32 bits to be divided differently for different areas of a network, which allows more nodes, networks, and subnetworks to be addressed, given a finite address space. [3]

Although, it addresses all of the problems with RIP, OSPF itself is not an absolutely perfect routing protocol (if there is such?). For small and medium-size networks, the basic services of OSPF work very well, enabling a wide range of robust TCP/IP and multiprotocol topologies. But in really large configurations, the huge number of router updates, that flow between routers, can become an issue. For instance, in a mesh network with over a hundred routers a single link change can precipitate a flood of thousands of link-state messages that propagate across the entire network from router to router. In each router, the database that stores these messages can grow to over a megabyte of live data. Each time there is a change to the network, routers must recalculate new routes. In very large OSPF networks, topology convergence can be delayed, while routers exchange link-state messages, update databases, and recalculate routes.

Areas

This delay in network convergence is the natural effect of very large topologies, and it will occur with any router protocol. Fortunately OSPF was designed and implemented in a much better way than RIP, to address this issue by what's known as the OSPF *area*. OSPF areas are simply logical subdivisions of an OSPF network. Typically, an enterprise is divided into areas that correspond to buildings, campuses, regions, or other administrative domains. An enterprise can have a practically unlimited number of areas.

OSPF routers within one area do not exchange topology updates with the routers in the other areas. When a LAN or WAN link is added to one area, topology updates flow only to routers within that area. This reduces the number of updates that flow through the network and the size of the topology databases in each router. In an enterprise with 500 routers, the creation of 10 areas of 50 routers means that each router only needs to store link information for 50 routers—not 500.

OSPF areas are connected to each other by means of a backbone that is just another area unto itself. A router that connects its area to the backbone must maintain a topology database for both areas. These special "multi-area" routers are called *area border routers*, and they serve as a filter for topology updates that move between areas and the backbone. Area border routers communicate with each other using special link-state messages that contain a short-hand summarization of the LAN/WAN topology in their areas.

continued on next page

Migration to OSPF is not a Luxury*(continued)*

Area border routers summarize topology information with an addressing technique that is analogous to U.S. Post Office zip codes. When a post office in New York City, for instance, handles a letter addressed to San Francisco, it only needs to look at the first few digits in the ZIP code to know that the letter is bound for SF; the remaining digits are not relevant at that point in the process. The national telephone system works the same way: when routing a call, major phone switches decode just the three-digit area-code prefix of a phone number to determine which long-distance trunk line to select; the rest of the digits are ignored until the call is routed to the right state and city.

These techniques of "hierarchical" addressing reduce the complexity of the phone and postal systems, because routing elements do not need to know all the details of end-to-end routes to every possible destination. Route information at the top of the hierarchy only relates to regions or areas. The OSPF area feature works the same way. Area border routers send-out special LSA update messages that advertise a range of IP addresses that reside in an area. Area border routers store these summarization messages in a special database that tells them how to forward inter-area traffic between areas. Route calculations based on summarization only need to determine what range an IP destination address falls within, based on the first few bytes of the address—not the complete address. However, this addressing feature does not come for free: before this process begins, OSPF summarization parameters must be administratively configured in area border routers. This task is similar to the configuration of router traffic filters or priorities.

A good example of OSPF areas' benefits can be seen in a campus environment, where each building is defined as an area. For example, le'ts take a campus where each building has 12 floors and a multiprotocol router on each floor. Without OSPF areas, routers would have to exchange updates with every other router on the campus, creating, in the process, topology databases that represent every routing node and link. When areas are deployed, routers only exchange link state information with routers in the same building. An area border router in each building forms a link between the building and the campus backbone.

Another OSPF area scenario is a national internetwork that is divided into areas that correspond to different regions of the country. For example, all the routers in the New York area would have identical databases that cover the New York region only, and the same would apply to other regions as well—Chicago, Tampa, etc. In each area, an area border router is attached to the national backbone. This approach eliminates the need to propagate router update messages across entire national (or international) internets. The area feature of OSPF is not the only factor in building large reliable networks, but it is one of the most important. When deploying OSPF areas, it is critical that area border routers have enough resources (memory, CPU) to handle topology chores for both the local and the backbone areas. Low-end routers may be overwhelmed by having to maintain multiple databases. Some additional considerations to keep in mind:

- Choose the correct topology (hub, mesh, hierarchical, collapsed backbone, etc.)
- Fine-tune OSPF timers within routers—the frequency of OSPF administrative messages often can be adjusted, and the right settings conserve bandwidth and speed convergence.

- Assign a meaningful OSPF "cost" to each link; OSPF computes its routes based on a least-cost metric; in configuring a network, each link is assigned a relative cost, based on its bandwidth, security, reliability, and so on—the ability of an OSPF network to adapt to end-user demands is directly proportional to the care invested in assigning cost values to links.

Summary

In many places, RIP is still used in TCP/IP networks, that have not been upgraded to OSPF. It is also used on OSPF networks as an end-station-to-router protocol.

RIP-2 provides very important enhancements to the original "ancient" RIP. However, it is not as widely supported as RIP-1. If both RIP-1 and RIP-2 are supported by specific products, the choice is simple: RIP-2 is certainly the one that has to be used in such environments. But, since it only resolves part of RIP's problems, OSPF should be the protocol of choice for intra-domain routing.

OSPF addresses all the deficiencies of RIP, without affecting connectivity to RIP based networks. Fast growing networks must be designed properly if the capabilities of OSPF are to be fully exploited. Because of its ability to handle variable networking masks OSPF also allows us not to waste today's so precious IP addresses.

Ideally, network design should include a consistent enterprise-wide IP address assignment policy that lends itself to the creation of OSPF areas and address summarization. If correct design and router-tuning takes place, OSPF will allow networks to scale to very large topologies, while maintaining high levels of availability and performance.

References

[1] Kleinrock, L., Farouk, K., "Hierarchical Routing for Large Networks," *Computer Networks,* Volume 1 (1977), North-Holland Publishing Company.

[2] *ConneXions,* Volume 3, No. 1, January 1989, Special Issue: Subnets.

[3] Fuller, V., Li, T., Yu, J., Varadhan, K., "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy," RFC 1519, September 1993.

[4] Malkin, G., "RIP Version 2 Carrying Additional Information," RFC 1723, November 1994.

[5] Moy, J., "OSPF Version 2," RFC 1583, March 1994.

[6] Hedrick, C., "Routing Information Protocol," RFC 1058, June 1988.

[7] *ConneXions,* Volume 3, No. 8, August 1989, Special Issue: Internet Routing.

[8] Tsuchiya, P. (now Francis, P.), "Components of OSI: IS–IS Intra-Domain Routing," *ConneXions,* Volume 3, No. 8, August 1989.

EDDIE RABINOVITCH has nineteen years of experience in Data Processing and Information Technology in general, and Data Communications in particular. Held senior technical and consulting positions with major corporation, including IBM, Dreyfus Corp., Dun & Bradstreet, and Chemical Bank. As of November 1995, Eddie joins the Network and Desktop Consulting Practice within the Global Services Support organization at Unisys Corporation as a Senior Manager. Eddie published more than 40 papers with several technical and trade publications, and currently has a regular monthly column "Net Perspectives" with *Client/Server Computing* magazine. Eddie is also serving as a member of the editorial review boards for the *Enterprise Systems Journal, Enterprise Internetworking Journal,* and The Computer Measurement Group. He is also a frequent speaker at national and international conferences. E-mail: `US000318@pop3.interramp.com`

Is Our Network Infrastructure Sound Enough?

by Peter G. Neumann, SRI International

Introduction

In the absence of extremely defensive system architectures, security and reliability are both weak-link phenomena. Unexpected problems can result from weak links being severed—whether accidentally or intentionally. Unfortunately, our operating systems are generally weak, our network protocols are vulnerable, and the networks themselves are risky (including the network nodes as well as unencrypted network media). Everything from the Internet Worm to the recent Netscape bugs reminds us repeatedly that major security flaws are rampant and can be exploited.

Historically, it is important for us to remember past problems and avoid their recurrence. However, people's memories of things past seem to be very short, prompting me to offer a few reminders. Two sets of similar cases immediately come to mind as particularly relevant to the *ConneXions* audience—computer network outages and telephone system outages.

Computer networks

The 1980 ARPANET case and the 1990 AT&T case are well known to long-time networkers, but bear retelling for younger folks as well as to remind the oldtimers that similar problems are still lurking. Both cases had similar origins, resulting from widespread propagation that began as a local problem.

ARPANET collapse

On October 27, 1980, the ARPANET experienced an unprecedented outage of approximately 4 hours, after years of almost flawless operation. This dramatic event is discussed in detail in a wonderful article by Eric Rosen ("Vulnerabilities of Network Control Protocols," ACM SIGSOFT *Software Engineering Notes,* 6, 1, 6–8, January 1981).

The collapse of the network resulted from an unforeseen interaction among three different problems: (1) a hardware failure resulted in bits being dropped in memory; (2) a redundant single-error-detecting code (simple parity checks) was used for transmission, but not for storage; and (3) the garbage-collection algorithm for removing old messages was not resistant to the simultaneous existence of one message with several different time stamps. This particular combination of circumstances had not arisen previously. In normal operation, each net node broadcasts a status message to each of its neighbors once per minute; 1 minute later, that message is then rebroadcast to the iterated neighbors, and so on. In the absence of bogus status messages, the garbage-collection algorithm is relatively sound. It keeps only the most recent of the status messages received from any given node, where *recency* is defined as the larger of two close-together 6-bit time stamps, modulo 64. Thus, for example, a node could delete any message that it had already received via a shorter path, or a message that it had originally sent that was routed back to it. For simplicity, 32 was considered a permissible difference, with the numerically larger time stamp being arbitrarily deemed the more recent in that case. In the situation that caused the collapse, the correct version of the time stamp was 44 [101100 in binary], whereas the bit-dropped versions had time stamps 40 [101000] and 8 [001000]. The garbage-collection algorithm noted that 44 was more recent than 40, which in turn was more recent than 8, which in turn was more recent than 44 (modulo 64). Thus, all three versions of that status message had to be kept.

From then on, the normal generation and forwarding of status messages from the particular node were such that *all of those messages* and their successors with newer time stamps had to be kept, thereby saturating the memory of each node.

In effect, this was a naturally propagating, globally contaminating effect. Ironically, the status messages had the highest priority, and thus defeated all efforts to maintain the network nodes remotely. Every node had to be shut down manually. Only after each site administrator reported back that the local nodes were down could the network be reconstituted; otherwise, the contaminating propagation would have begun anew.

AT&T system runaway

In mid-December 1989, AT&T installed new software in 114 electronic switching systems (Number 4 ESS), intended to reduce the overhead required in signaling between switches by eliminating a signal indicating that a node was ready to resume receiving traffic; instead, the other nodes were expected to recognize implicitly the readiness of the previously failed node, based on its resumption of activity. Unfortunately, there was an undetected latent flaw in the recovery-recognition software in every one of those switches.

On January 15, 1990, one of the switches experienced abnormal behavior; it signaled that it could not accept further traffic, went through its recovery cycle, and then resumed sending traffic. A second switch accepted the message from the first switch and attempted to reset itself. However, a second message arrived from the first switch that could not be processed properly, because of the flaw in the software. The second switch shut itself down, recovered, and resumed sending traffic. That resulted in the same problem propagating to the neighboring switches, and then iteratively and repeatedly to all 114 switches. The hitherto undetected problem manifested itself in subsequent simulations whenever a second message arrived within too short a time. AT&T finally was able to diagnose the problem and to eliminate it by reducing the messaging load of the network, after a 9-hour nationwide blockade. With the reduced load, the erratic behavior effectively went away by itself, although the software still had to be patched correctly to prevent a recurrence. Reportedly, approximately 5 million calls were blocked.

The ultimate cause of the problem was traced to a **C** program that contained a `break` statement within an `if` clause nested within a `switch` clause. This problem can be called a programming error, or a deficiency of the **C** language and its compiler, depending on your taste, in that the intervening `if` clause was in violation of expected programming practice.

Telephone systems

Numerous recent cable cuts in telephone lines (shutting down airports and communications infrastructure generally) remind us of the partitioning of the ARPANET December 12, 1986, in which New England was separated from the rest of the net, due to a single-point failure—seven alternative lines all went through the same conduit in White Plains, New York, and all seven were accidentally severed. Did anyone learn from that episode? Oh, yes. The Associated Press insisted on having two different cables. And yet, those two cables were adjacent to one another, so that a single backhoe in Annandale, Virginia, could take out both cables in a single swipe on June 14, 1991. Recent outages also occurred in Chicago (November 19, 1990), Newark, New Jersey (January 4, 1991), and San Francisco (July 15, 1991). Other serious outages affected Washington, D.C., Los Angeles, and Pittsburgh (all on June 27, 1991, due to a faulty software patch in the *Signaling System No. 7* (SS7) switching systems), and New York City (September 17, 1991, due to operational problems that caused the system to be disconnected from its main power, the backup generator hookup to fail, and the standby batteries to be drained).

continued on next page

Network Infrastructure Sound Enough? *(continued)*

Conclusions

The 1980 ARPANET outage was triggered by hardware failures, but required the presence of a software weakness and a hardware design shortcut that permitted total propagation of the contaminating effect. The 1990 AT&T blockage resulted from a programming mistake and a network design that permitted total propagation of the contaminating effect. The 1986 ARPANET separation was triggered by an environmental accident, but depended on a poor implementation decision. The flurry of cable cuttings resulted from digging accidents, but the effects in each case were exacerbated by designs that made the systems particularly vulnerable. The SS7 problems resulted from a faulty code patch and the absence of testing before that patch was installed. The New York City telephone outage resulted from poor network administration, and was complicated by the earlier removal of warning alarms.

Further details on the risks in communications and in our computer infrastructure can be found in *Computer-Related Risks,* by Peter G. Neumann, Addison-Wesley, 1995, ISBN 0-201-55805-X, from which some of the above material is excerpted. That book also includes some analysis of how accidental problems could alternatively been triggered maliciously, and how maliciously caused problems could alternatively have been triggered accidentally.

In summary, the primary causes of the cited communications problems were largely environmental, or were the results of problems in maintenance and system evolution. Hardware malfunctions were involved in most cases, but generally as secondary causative factors. Software was involved in several cases. Many of the cases have multiple causative factors, and originate in bad design and inadequate implementation.

All of these cases suggest that we must work much harder to make our networking infrastructure more resilient—more reliable, and more secure. The existing infrastructure is not ready for prime time.

PETER G. NEUMANN received AB, SM, and PhD degrees from Harvard in 1954, 1955, 1961, respectively, and a Dr rerum naturarum from the Technische Hochschule, Darmstadt, Germany, in 1960. He has worked in the computer field since 1953. In the Computer Science Lab at Bell Telephone Labs at Murray Hill, New Jersey throughout the 1960s, he was involved in research in computers and communications; during 1965–69, he participated in the design and development of Multics, jointly with MIT and Honeywell. He was a visiting Mackay Lecturer at Stanford in 1964 and Berkeley in 1970–71. He has been in the Computer Science Laboratory at SRI since 1971, where he has been concerned with computer systems having requirements for security, reliability, human safety, and high assurance (including formal methods). He participated in the 1981–82 Air Force Studies Board database security study, and the 1989–90 National Research Council System Security Study Committee that produced the report, *Computers at Risk.* He also served on an Expert Panel for the U.S. House Judiciary Committee Subcommittee on Civil and Constitutional Rights. He recently served on an ACM panel ("Codes, Keys, and Conflicts: Issues in U.S. Crypto Policy," ACM, June 1994), and is now on a new National Research Council study group, both of which are targeted at reviewing U.S. crypto policy. For the Association for Computing Machinery, he was founder and Editor of the SIGSOFT *Software Engineering Notes* (1976–1993) and is now Associate Editor for the RISKS material; Chairman of the ACM Committee on Computers and Public Policy (since 1985); and a Contributing Editor for CACM (since 1990) for the monthly "Inside Risks" column. In 1985 he created, and still moderates, the ACM Forum on Risks to the Public in the Use of Computers and Related Technology, which is one of the most widely read of the on-line computer newsgroups. RISKS (`comp.risks`) provides a medium for discussion of issues relating to all aspects of computers and the social and technological problems that they create. Neumann is a Fellow of both the ACM and the Institute of Electrical and Electronics Engineers. E-mail: `Neumann@CSL.SRI.COM`

Future NetWorld+Interop Dates and Locations

NetWorld+Interop 96	Las Vegas, NV	April 1–5, 1996
NetWorld+Interop 96	Frankfurt, Germany	June 10–14, 1996
NetWorld+Interop 96	Tokyo, Japan	July 22–26, 1996
NetWorld+Interop 96	Atlanta, GA	September 16–20, 1996
NetWorld+Interop 96	Paris, France	October 7–11, 1996
NetWorld+Interop 96	London, England	Oct. 28–Nov. 1, 1996
NetWorld+Interop 96	Sydney, Australia	November 25–29, 1996

All dates are subject to change.

More information

Call 1-800-INTEROP or +1-415-578-6900 for more information. Or send e-mail to `info@interop.com` or fax to +1-415-525-0194. For the latest information about NetWorld+Interop including *N+I Online!* as well as other SOFTBANK produced events, check our home page at `http://www.interop.com`

NetWorld+Interop is produced by SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California 94404–1138, USA.

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City
California 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`
URL: `http://www.interop.com`

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our new subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

Call for Papers

INET '96, the 6th Annual Conference of the Internet Society focusing on worldwide issues of Internet networking will be held 25–28 June 1996 in Montreal, Canada. This conference brings together those extending the reach and use of Internet networks. Participants include those developing and implementing Internet networks, applications, and policies for worldwide infrastructure development. The development of Internet networks in an ever wider variety of social, cultural, economic and linguistic contexts is also a focal point of this conference. INET '96 will encompass certain horizontal threads reflecting the general tone of this conference. In particular, the desire to treat the Internet as a unified, complex, phenomenon meshing highly technical issues with deeply social, economic, and cultural concerns is stressed in order to help the whole world better understand the Internet revolution.

Topics

Topics for submissions include but are not limited to the following:

• *Internet Applications and Services:* The Internet provides a foundation for the delivery of many advanced services. The technologies to deliver these services include advanced tools for managing, searching, and accessing distributed information. They also include techniques for dealing with multimedia, files systems, computing, collaboration, user interfaces, multiple language support and mobility.

• *Transforming Internet Commerce and Reshaping the Marketplace:* The Internet and its related technologies provide an important platform for transformation of business and commercial activities. Business activities continue to evolve on the Internet. New product offerings such as commerce servers, publishing servers, community servers and electronic malls have captured the imagination of the public and many business leaders. Internet networks deeply transform the reach of firms, allowing small companies to have global reach. New forms of competition emerge with related questions about the nature and security of transactions, the need for new electronic currencies. New customer relationships emerge with implications for advertising and distribution and delivery of products and services.

• *Internet Learning and Teaching:* The Internet provides unparalleled richness from the standpoint of the individual learner. Focused attention on organization and presentation of teaching and learning material in a highly interactive environment produces new learning and teaching paradigms. Organizations of all kinds including primary and secondary schools, post-secondary education institutions, government institutions and commercial enterprises seek to use the Internet and its related technologies to enhance the learning and teaching process. The application of Internet technologies to education accelerates such developments as "just-in-time learning." Some of these trends deeply reshape functions and objectives of traditional learning institutions. Experiments with new teaching applications and the building of global communities also transform the nature of education.

• *Networking Technology Frontiers:* The increasing sophistication of network applications and enormous growth in number of people using the Internet demand new networking solutions. Advanced technologies and services to expand, rationalize and manage core network services develop quickly. Networking designs, protocols, registry processes and services, transport services and security requirements continue to undergo rapid evolution to meet the growing demand.

• *Internet and Social Transformations:* The global Internet is affecting how people interact and how society works. Ideas and opinions flow faster and in new directions, and as a result power is being distributed in unexpected ways. Until the Internet, the growth of mass media pointed to a world with an increasingly homogeneous culture. Now, the Internet holds the promise to enhance cultural and linguistic diversity on a global scale. New kinds of communities are coming to light. Borders become porous to ideas, opinions, rumors and facts. Politics and governments are changing. If the Internet is truly the equivalent of printing with movable type, what can we already say about its effect on our societies?

• *Growing and Regulating the Internet:* Economic and Policy Issues More countries and the international community recognize Internet evolution as an important economic and policy issue. Major challenges continue as global and national communities struggle to understand the incremental nature of Internet evolution and how to encourage, regulate or discourage its use and growth. Advancing Internet technologies also cause redefinition of current economic activities, regulatory and economic policies, and political issues.

• *Expanding and Enhancing Internet Access:* Most parts of the world struggle to provide reliable access with reasonable performance. Many geographic areas also struggle to extend access to more individuals and institutions. Technical, economic, social and political barriers and solutions continue to evolve. Projects within geographic regions, countries and industries illustrate the nature of the challenges and the dimensions of potential solutions.

• *Internet Case Studies:* Individuals, organizations and governments use the Internet for a wide range of activities. These experiences, both successes and failures, form an important knowledge base of information about the Internet and also help define frontiers for further exploration and development.

Submissions

The official language of the conference is English. Papers will be selected based on full papers. Each submission must contain a separate one-page abstract with the title or topic, the names of the author(s), organizational affiliation(s), addresses, telephone number, fax number, and e-mail addresses and must identify a single point of contact if more than one author is listed. Abstracts should also include a keyword list, tied to the topics listed above. Upon acceptance papers must be resubmitted in the format required for publication in the proceedings. Papers in plain ASCII text should be submitted by *January 15, 1996* to: `inet-submission@isoc.org`. The INET '96 Program Committee can be contacted at: `inet-program@isoc.org`.

Workshops

INET '96 will be preceded by a seven-day *Developing Country Workshop* featuring intensive instruction with a hands-on emphasis on Internet set up, operations, maintenance and management. For more information, please send e-mail to: `workshop-info@isoc.org`.

INET '96 will also be preceded by a tentative two day program bringing together active Kindergarten through Secondary School Internet innovators from around the world to share experiences and learn new advanced tools and collaboration techniques. For information about the *Primary and Secondary School Workshop,* please send e-mail to: `inet-k12@isoc.org`.

More information

URL: `http://www.isoc.org/conferences/inet96/`
E-mail: `inet96@isoc.org`
Tel: +1 703 648 9888 • Fax: +1 703 648 9887

CONNEXIONS
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

Printed on recycled paper